arXiv:1607.07549v2 [math.ST] 11 Sep 2016

Concentration of Measure for Radial Distributions and Consequences for Statistical Modeling

Ery Arias-Castro Xiao Pu

Abstract

Motivated by problems in high-dimensional statistics such as mixture modeling for classification and clustering, we consider the behavior of radial densities as the dimension increases. We establish a form of concentration of measure, and even a convergence in distribution, under additional assumptions. This extends the well-known behavior of the normal distribution (its concentration around the sphere of radius square-root of the dimension) to other radial densities. We draw some possible consequences for statistical modeling in high-dimensions, including a possible universality property of Gaussian mixtures.

1 Introduction

Nonparametric density estimation can be quickly difficult in high-dimensions because of the curse of dimensionality. Additional assumptions are often needed. The most popular one might well be the Naive Bayes approach, popular in classification (Lewis, 1998), which presumes that the variables are independent, or equivalently, that the density is the product of its marginals.

Another possibility is to assume that the density is elliptical, a classical assumption in multivariate analysis (Anderson, 2003), meaning that f is of the form $f(x) = |A|g(\|Ax\|)$, where A is a positive definite matrix. Any centered and non-degenerate normal distribution is elliptical, with base distribution the standard normal distribution and $A = \Sigma^{-1/2}$, where Σ is the covariance matrix. Note that the same distribution is also the product of its marginals when Σ is diagonal — the assumption underlying Linear Discriminant Analysis, an important parametric special case of Naive Bayes. High-dimensional density estimation is of course crucial in classification and clustering.

Our motivation comes from such problems, and in particular, mixture modeling in high-dimensions using elliptical distributions. In the process of working on this problem we uncovered a difficulty which we elaborate upon in the present paper: that of estimating the base function. Indeed, it is quite tempting to extend the Gaussian mixture models to models of the form

$$\sum_{k=1}^{K} \pi_k |A_k| g_k(\|A_k x\|), \tag{1}$$

where we assume the mixture has K components, with the kth component having weight π_k and density $|A_k|g_k(\|A_k x\|)$. For example, Bickel et al. (1998) and more recently Bhattacharyya and Bickel (2015) consider models of this kind. Instead of smoothness assumptions, we are more interested here in shape assumptions, for example that g_k is decreasing and/or log-concave on $\mathbb{R}_+$.

Both authors are with the Department of Mathematics, University of California, San Diego, USA. Contact Ery Arias-Castro or Xiao (Victor) Pu. We are grateful to Chris Sherlock and Daniel Elton for bringing their paper to our attention, and to Gabor Lugosi for helpful discussions. This work was partially supported by a grant from the US Office of Naval Research (N00014-13-1-0257).

Chang and Walther (2007) consider such mixture models but under the Naive Bayes assumption instead of assuming the densities are elliptical. An EM approach to fitting such a model involves being able to estimate g_k based on a sample from $g_k(\|x\|)$. And this is what we found challenging in our investigation.

Focusing on this task, suppose we have an IID sample from f, where f rotationally invariant (aka radial), meaning that $f(x) = g(\|x\|)$ for some function g, and consider the problem of estimating g. In fact, we can work with the magnitudes (the norms of the observations), which are sufficient. We explain the difficulty of estimating g by the fact that the magnitudes are highly concentrated as the dimension becomes large.

The simplest case of this concentration of measure phenomenon arises when we assume that that $g \propto \psi$, where ψ is fixed, as is the case in the Gaussian setting. Specifically, we assume we are in dimension $d+1$ and we work with ψ satisfying the following assumptions (and some additional assumptions specified later on)

$$\psi : \mathbb{R}_+ \to \mathbb{R}_+ \text{ is such that } 1/c_d := \int_0^\infty u^d \psi(u) \mathrm{d}u < \infty \text{ for all } d \geq 1. \tag{2}$$

We let X_d denote a random variable with density $c_d\psi(\|x\|)$ on $\mathbb{R}^{d+1}$ and let U_d denote its magnitude, $U_d = \|X_d\|$, which has density $c_d u^d \psi(u)$ on $\mathbb{R}_+$.

In this context we show a form of concentration of measure, and convergence in distribution, as the dimension d increases. Concentration is a well-known phenomenon in high-dimensions, in particular for product distributions (Naive Bayes), with far-reaching consequences (Boucheron et al., 2013; Ledoux, 2005). For radial distributions, it is not as well-known, except for when the density is Gaussian or uniform on a ball. (The latter is often used to explain some forms of curse of dimensionality.) This case was recently studied in detail by Sherlock and Elton (2012), who cite older work by Diaconis and Freedman (1984, 1987) in rather special cases. We comment on (Sherlock and Elton, 2012) in more detail in Section 4, after we present our results.

In the remaining of the paper we study what happens when $d \to \infty$. In Section 2 we study the case where ψ has compact support, which is the simplest situation. In Section 3 we consider the case where ψ is *not* compactly supported. In Section 4 we discuss our results and some possible implications for statistical modeling.

2 The case of compact support

In this whole section we assume that ψ has compact support. Define the supremum of the support as follows

$$u_* = \sup\Big\{u : \int_{u-\varepsilon}^{u} \psi(u)\mathrm{d}u > 0 \text{ for all } \varepsilon > 0\Big\}. \tag{3}$$

Note that $u_* < \infty$ by assumption and that the support of ψ is included in $[0, u_*]$. If ψ is continuous (which the reader can assume without much loss of generality), then the following is an equivalent definition $u_* = \sup\{u : \psi(u) > 0\}$. The emblematic example is that of the uniform distribution on the unit ball, in which case $\psi(u) = \mathbb{I}\{u \leq 1\}$ and $u_* = 1$. This distribution is well-known to concentrate near the boundary of its support (the unit sphere). Our results below extend this to other distributions with compact support.

2.1 Convergence in probability

We start by establishing a convergence in probability.

Theorem 1. *In the setting considered here, $U_d \to u_*$ in probability as $d \to \infty$.*

Proof. Assume $u_* = 1$ without loss of generality. Then

$$\mathbb{P}(U_d < 1-\varepsilon) = c_d \int_0^{1-\varepsilon} u^d \psi(u)\mathrm{d}u \le c_d(1-\varepsilon)^d \int_0^1 \psi(u)\mathrm{d}u, \tag{4}$$

while

$$\mathbb{P}(U_d \ge 1-\varepsilon) \ge \mathbb{P}(U_d \ge 1-\varepsilon/2) = c_d \int_{1-\varepsilon/2}^1 u^d \psi(u)\mathrm{d}u \ge c_d(1-\varepsilon/2)^d \int_{1-\varepsilon/2}^1 \psi(u)\mathrm{d}u. \tag{5}$$

Note that the last integral is strictly positive for all $\varepsilon > 0$ by definition of u_* in (3) (recall that we assumed that $u_* = 1$). Hence

$$\frac{\mathbb{P}(U_d < 1-\varepsilon)}{\mathbb{P}(U_d \ge 1-\varepsilon)} \le \frac{(1-\varepsilon)^d \int_0^1 \psi(u)\mathrm{d}u}{(1-\varepsilon/2)^d \int_{1-\varepsilon/2}^1 \psi(u)\mathrm{d}u} \to 0, \quad d \to \infty, \tag{6}$$

when $\varepsilon \in (0,1)$ is fixed. Since $\mathbb{P}(U_d < 1-\varepsilon) + \mathbb{P}(U_d \ge 1-\varepsilon) = 1$, we proved that $\mathbb{P}(U_d < 1-\varepsilon) \to 0$ for all $\varepsilon > 0$. This, coupled with the fact that $\mathbb{P}(U_d \le 1) = 1$, proves that $U_d \to 1$ in probability as $d \to \infty$. □

2.2 Convergence in distribution

Beyond a convergence in probability, we can establish a convergence in distribution. The limiting distribution happens to depend on the behavior of ψ in the neighborhood of u_*. We only cover the case where ψ behaves as a power function near u_*.

Theorem 2. *In the setting considered here, assume in addition that ψ is bounded and that $\psi(u) \sim a(u_* - u)^b$ as $u \nearrow u_*$ for some $a > 0$ and $b > -1$. Then $d(u_* - U_d)$ converges weakly to the Gamma distribution with shape parameter $b+1$ and rate $1/u_*$.*

Proof. Assume without loss of generality that $u_* = 1$. We first control the behavior of c_d as $d \to \infty$. Fix $\varepsilon \in (0,1)$. On the one hand, by the assumptions on ψ and Dominated Convergence, we have

$$\int_{1-\varepsilon}^1 u^d \psi(u)\mathrm{d}u \sim \int_{1-\varepsilon}^1 u^d a(1-u)^b \mathrm{d}u \sim aB(d+1, b+1), \quad d \to \infty, \tag{7}$$

where B is the Beta function. On the other hand, by Theorem 1,

$$c_d \int_{1-\varepsilon}^1 u^d \psi(u)\mathrm{d}u \sim 1, \quad d \to \infty. \tag{8}$$

Together, this proves that

$$1/c_d \sim aB(d+1, b+1) \sim a\Gamma(b+1)d^{-(b+1)}, \quad d \to \infty, \tag{9}$$

where Γ is the Gamma function.

We now consider the case where $\varepsilon = \varepsilon_d \to 0$ as $d \to \infty$. More precisely, we fix $t > 0$ and set $\varepsilon_d = t/d$. By Dominated Convergence again, applied twice, and a change of variables, as $d \to \infty$, we have

$$\begin{aligned}
\mathbb{P}(U_d > 1 - t/d) &= c_d \int_{1-t/d}^1 u^d \psi(u)\mathrm{d}u && (10)\\
&\sim c_d \int_{1-t/d}^1 u^d a(1-u)^b \mathrm{d}u && (11)\\
&= c_d a d^{-(b+1)} \int_0^t (1 - v/d)^d v^b \mathrm{d}v && (12)\\
&\sim \frac{1}{\Gamma(b+1)} \int_0^t e^{-v} v^b \mathrm{d}v. && (13)
\end{aligned}$$

Recognizing the distribution function of the Gamma distribution with shape parameter $b+1$ and rate 1, the proof is complete. □

3 The case of non-compact support

We now assume that ψ has non-compact support, which is equivalent to $u_* = \infty$ in (3). We note that here the emblematic example is that of the standard normal distribution, which is known to concentrate near the sphere of radius $\sqrt{d}$, meaning $U_d/\sqrt{d} \to 1$. In fact, U_d^2 has the chi-squared distribution with d degrees of freedom, and in particular, $\sqrt{2}(U_d - \sqrt{d})$ is asymptotically standard normal in the limit $d \to \infty$. Our results below extend this phenomena to other radial distributions.

While we were able to handle the case of compact support, which we treated in Section 2, with very natural assumptions, the case of non-compact support appears more challenging and our working assumptions are more complicated. This is despite the fact that we favored simplicity over generality. Nevertheless, our working assumptions include interesting (and natural) examples.

3.1 Convergence in probability

We start by establishing a convergence in probability.

We start by making the following assumptions. We assume there is $u_\ddagger$ such that, for $u \geq u_\ddagger$, $\Lambda(u) := -\log \psi(u)$ is differentiable and $L(u) := u\Lambda'(u)$ is increasing. In addition, we assume that $M(u) := L(u)/\log(u) \to \infty$ as $u \to \infty$ and

$$\limsup_{u\to\infty} \frac{M((1-\varepsilon)u)}{M(u)} \leq 1, \quad \liminf_{u\to\infty} \frac{M((1+\varepsilon)u)}{M(u)} \geq 1, \quad \forall \varepsilon \in (0,1). \tag{14}$$

Theorem 3. *In the setting considered here, $U_d/u_d \to 1$ in probability as $d \to \infty$, where $u_d := L^{-1}(d)$.*

Example 1. Consider the case where $\Lambda(u) = c\log(u+a)^\alpha (u+b)^\beta$, where $a > 0$, $b \geq 0$, $c > 0$, $\alpha \in \mathbb{R}$ and $\beta > 0$. Surely, this defines a bonafide shape function ψ in the sense of (2). It can be shown that ψ defined as such satisfies the conditions of Theorem 3, with

$$u_d \sim c^{-1/\beta}\beta^{(\alpha-1)/\beta}(\log d)^{-\alpha/\beta} d^{1/\beta}, \quad d \to \infty. \tag{15}$$

Proof. The function $u \mapsto u^d\psi(u)$ is increasing on $[0, u_d)$ and decreasing on (u_d, ∞). Indeed, $\log(u^d\psi(u)) = d\log u - \Lambda(u)$ has derivative $\frac{1}{u}(d - L(u))$, which is positive for $u < u_d$, zero at $u = u_d$, and negative at $u > u_d$, by our assumptions and the definition of u_d. Note that, necessarily, $u_d \to \infty$ as $d \to \infty$.

We have

$$\mathbb{P}(U_d \leq v) = c_d \int_0^v u^d \psi(u) \mathrm{d}u. \tag{16}$$

Fix $\varepsilon \in (0,1)$.

Left tail. Using the fact that $u^d\psi(u) \leq u_0^d\psi(u_0)$ for any $u \leq u_0 \leq u_d$, we have

$$\frac{1}{c_d}\mathbb{P}(U_d \leq (1-\varepsilon)u_d) = \int_0^{(1-\varepsilon)u_d} u^d\psi(u)\mathrm{d}u \tag{17}$$

$$\leq ((1-\varepsilon)u_d)^{d+1}\psi((1-\varepsilon)u_d), \tag{18}$$

and we also have

$$\frac{1}{c_d}\,\mathbb{P}(U_d \geq (1-\varepsilon)u_d) = \int_{(1-\varepsilon)u_d}^{\infty} u^d \psi(u)\mathrm{d}u \tag{19}$$

$$\geq \int_{(1-\varepsilon/2)u_d}^{u_d} u^d \psi(u)\mathrm{d}u \tag{20}$$

$$\geq (\varepsilon/2)u_d((1-\varepsilon/2)u_d)^d \psi((1-\varepsilon/2)u_d). \tag{21}$$

Taking the ratio, we obtain

$$\frac{\mathbb{P}(U_d \leq (1-\varepsilon)u_d)}{\mathbb{P}(U_d \geq (1-\varepsilon)u_d)} \leq \frac{((1-\varepsilon)u_d)^{d+1}\psi((1-\varepsilon)u_d)}{(\varepsilon/2)u_d((1-\varepsilon/2)u_d)^d \psi((1-\varepsilon/2)u_d)} \tag{22}$$

$$\leq \frac{1-\varepsilon}{\varepsilon/2}\left(\frac{1-\varepsilon}{1-\varepsilon/2}\right)^d \frac{\psi((1-\varepsilon)u_d)}{\psi((1-\varepsilon/2)u_d)}. \tag{23}$$

Applying the logarithm, and ignoring the constant factor, we further get

$$d\log\left(\frac{1-\varepsilon}{1-\varepsilon/2}\right) - \Lambda((1-\varepsilon)u_d) + \Lambda((1-\varepsilon/2)u_d) \tag{24}$$

$$= -d\int_{(1-\varepsilon)u_d}^{(1-\varepsilon/2)u_d} \frac{1}{u}\mathrm{d}u + \int_{(1-\varepsilon)u_d}^{(1-\varepsilon/2)u_d} \Lambda'(u)\mathrm{d}u \tag{25}$$

$$= -\int_{(1-\varepsilon)u_d}^{(1-\varepsilon/2)u_d} (d - L(u))\frac{1}{u}\mathrm{d}u \tag{26}$$

$$\leq -(d - L((1-\varepsilon/2)u_d))\log\left(\frac{1-\varepsilon/2}{1-\varepsilon}\right), \tag{27}$$

where we used the monotonicity of L in the last line. Therefore, to show that the fraction in (22) converges to 0, it suffices to show that $d - L((1-\varepsilon/2)u_d) \to \infty$. The limit is as $d \to \infty$ while ε remains fixed. Using the fact that $L(u_d) = M(u_d)\log u_d = d$, we have

$$d - L((1-\varepsilon/2)u_d) = d - M((1-\varepsilon/2)u_d)\log((1-\varepsilon/2)u_d) \tag{28}$$

$$= d - M((1-\varepsilon/2)u_d)\log(1-\varepsilon/2) - M((1-\varepsilon/2)u_d)\frac{d}{M(u_d)} \tag{29}$$

$$= -M((1-\varepsilon/2)u_d)\log(1-\varepsilon/2) + d\left[1 - \frac{M((1-\varepsilon/2)u_d)}{M(u_d)}\right]. \tag{30}$$

In the last line, the first term tends to infinity because $M(u) \to \infty$ as $u \to \infty$ and $u_d \to \infty$, while the second terms is nonnegative in the limit because of (14), so that the last expression tends to infinity.

We conclude that, for the left tail,

$$\mathbb{P}(U_d \leq (1-\varepsilon)u_d) \to 0, \quad d \to \infty. \tag{31}$$

Right tail. Using the fact that $u^\ell\psi(u) \leq u_0^\ell\psi(u_0)$ for any $u \geq u_0 \geq u_\ell$, where $u_\ell = L^{-1}(\ell)$ in

congruence with our definition above, and assuming for now that $b := L((1+\varepsilon)u_d) > d+1$, we have

$$\frac{1}{c_d}\,\mathbb{P}(U_d \geq (1+\varepsilon)u_d) = \int_{(1+\varepsilon)u_d}^{\infty} u^d \psi(u)\mathrm{d}u \tag{32}$$

$$\leq ((1+\varepsilon)u_d)^b \psi((1+\varepsilon)u_d) \int_{(1+\varepsilon)u_d}^{\infty} u^{d-b}\mathrm{d}u \tag{33}$$

$$= ((1+\varepsilon)u_d)^b \psi((1+\varepsilon)u_d) \frac{((1+\varepsilon)u_d)^{d+1-b}}{b-d-1} \tag{34}$$

$$= \frac{((1+\varepsilon)u_d)^{d+1}}{b-d-1}\psi((1+\varepsilon)u_d), \tag{35}$$

and we also have

$$\frac{1}{c_d}\,\mathbb{P}(U_d \leq (1+\varepsilon)u_d) = \int_{0}^{(1+\varepsilon)u_d} u^d \psi(u)\mathrm{d}u \tag{36}$$

$$\geq \int_{u_d}^{(1+\varepsilon/2)u_d} u^d \psi(u)\mathrm{d}u \tag{37}$$

$$\geq (\varepsilon/2)u_d((1+\varepsilon/2)u_d)^d \psi((1+\varepsilon/2)u_d). \tag{38}$$

Taking the ratio, we obtain

$$\frac{\mathbb{P}(U_d \geq (1+\varepsilon)u_d)}{\mathbb{P}(U_d \leq (1+\varepsilon)u_d)} \leq \frac{\frac{((1+\varepsilon)u_d)^{d+1}}{b-d-1}\psi((1+\varepsilon)u_d)}{(\varepsilon/2)u_d((1+\varepsilon/2)u_d)^d \psi((1+\varepsilon/2)u_d)} \tag{39}$$

$$\leq \frac{1+\varepsilon}{\varepsilon/2}\frac{1}{b-d-1}\left(\frac{1+\varepsilon}{1+\varepsilon/2}\right)^d \frac{\psi((1+\varepsilon)u_d)}{\psi((1+\varepsilon/2)u_d)}. \tag{40}$$

We pause to show that $b-d \to \infty$ eventually. This is because, using the fact that $L(u_d) = M(u_d)\log u_d = d$,

$$b-d = L((1+\varepsilon)u_d) - d = M((1+\varepsilon)u_d)\log((1+\varepsilon)u_d) - d \tag{41}$$

$$= M((1+\varepsilon)u_d)\log(1+\varepsilon) + \left[\frac{M((1+\varepsilon)u_d)}{M(u_d)} - 1\right]d. \tag{42}$$

In the last line, the first term tends to infinity because $M(u) \to \infty$ as $u \to \infty$ and $u_d \to \infty$, while the second terms is nonnegative in the limit because of (14), so that the last expression tends to infinity.

Returning to (39), applying the logarithm, and ignoring the first two factors whose product is bounded by 1 eventually, we further get

$$d\log\left(\frac{1+\varepsilon}{1+\varepsilon/2}\right) - \Lambda((1+\varepsilon)u_d) + \Lambda((1+\varepsilon/2)u_d) \tag{43}$$

$$= d\int_{(1+\varepsilon/2)u_d}^{(1+\varepsilon)u_d} \frac{1}{u}\mathrm{d}u - \int_{(1+\varepsilon/2)u_d}^{(1+\varepsilon)u_d} \Lambda'(u)\mathrm{d}u \tag{44}$$

$$= -\int_{(1+\varepsilon/2)u_d}^{(1+\varepsilon)u_d} (L(u)-d)\frac{1}{u}\mathrm{d}u \tag{45}$$

$$\leq -(L((1+\varepsilon/2)u_d) - d)\log\left(\frac{1+\varepsilon}{1+\varepsilon/2}\right), \tag{46}$$

where we used the monotonicity of L in the last line. Therefore, to show that the fraction in (39) converges to 0, it suffices to show that $L((1+\varepsilon/2)u_d) - d \to \infty$, and we already did this in (41).

We conclude that, for the right tail,

$$\mathbb{P}(U_d \geq (1+\varepsilon)u_d) \to 0, \quad d \to \infty. \tag{47}$$

We can therefore conclude that $U_d/u_d \to 1$ in probability as $d \to \infty$. □

3.2 Convergence in distribution

We now turn to establishing a convergence in distribution. Although we speculate that other cases may arise, we give (additional) sufficient conditions for a Gaussian limit.

We make the following additional assumptions. We assume that L is differentiable with $\nu_d := u_d L'(u_d) \to \infty$ and that there is $\omega_d \to \infty$ such that

$$\left|L(u_d) - L((1-\varepsilon)u_d) - \varepsilon u_d L'(u_d)\right| \leq |\varepsilon|\nu_d/\omega_d, \quad \text{whenever } \varepsilon^2 \leq \omega_d/\nu_d. \tag{48}$$

Note that (48) is a form of first-order Taylor expansion around u_d.

The following refines Theorem 3.

Proposition 1. *In the setting considered here,*

$$\mathbb{P}(1-\varepsilon_d \leq U/u_d \leq 1+\varepsilon_d) \to 1, \quad \textit{whenever } \varepsilon_d \gg 1/\sqrt{\nu_d}. \tag{49}$$

Proof. By mononicity in $\varepsilon_d > 0$, it is enough to show that when $\varepsilon_d^2 \leq \omega_d/\nu_d$. Then to prove (49), as before, it suffices to show that

$$\text{(left tail) } \frac{\mathbb{P}(U_d \leq (1-\varepsilon_d)u_d)}{\mathbb{P}(U_d \geq (1-\varepsilon_d)u_d)} \to 0 \quad \text{and} \quad \text{(right tail) } \frac{\mathbb{P}(U_d \geq (1+\varepsilon_d)u_d)}{\mathbb{P}(U_d \leq (1+\varepsilon_d)u_d)} \to 0. \tag{50}$$

Left tail. As before, we can show that $b := L((1-\varepsilon_d)u_d)$ satisfies $b - d \to -\infty$, so that we may assume that $b < d-1$. Then using the fact that $u^\ell \psi(u) \leq u_0^\ell \psi(u_0)$ for any $u \leq u_0 \leq u_\ell$, where $u_\ell = L^{-1}(\ell)$, we have

$$\frac{1}{c_d}\mathbb{P}(U_d \leq (1-\varepsilon_d)u_d) = \int_0^{(1-\varepsilon_d)u_d} u^d \psi(u)\mathrm{d}u \tag{51}$$

$$\leq ((1-\varepsilon_d)u_d)^b \psi((1-\varepsilon_d)) \int_0^{(1-\varepsilon_d)u_d} u^{d-b}\mathrm{d}u \tag{52}$$

$$= \frac{((1-\varepsilon_d)u_d)^{d+1}}{d-b+1}\psi((1-\varepsilon_d)u_d). \tag{53}$$

Taking the ratio of (53) to (21) (but replacing ε by ε_d), we obtain

$$\frac{\mathbb{P}(U_d \leq (1-\varepsilon_d)u_d)}{\mathbb{P}(U_d \geq (1-\varepsilon_d)u_d)} \leq \frac{1}{d-b+1}\frac{1-\varepsilon_d}{\varepsilon_d/2}\left(\frac{1-\varepsilon_d}{1-\varepsilon_d/2}\right)^d \frac{\psi((1-\varepsilon_d)u_d)}{\psi((1-\varepsilon_d/2)u_d)}. \tag{54}$$

As in (23) and (27), we apply a logarithm, and obtain the upper bound

$$\log\frac{1-\varepsilon_d}{\frac{\varepsilon_d}{2}(d-b+1)} - [d - L((1-\varepsilon_d/2)u_d)]\log\frac{1-\varepsilon_d/2}{1-\varepsilon_d}. \tag{55}$$

By (48), which is applicable by our assumption $\varepsilon_d^2 \le \omega_d/\nu_d$, we have

$$d - b = L(u_d) - L((1-\varepsilon_d)u_d) = \varepsilon_d\nu_d \pm |\varepsilon_d|\nu_d/\omega_d \sim \varepsilon_d\nu_d, \quad d \to \infty. \tag{56}$$

Similarly,

$$d - L((1-\varepsilon_d/2)u_d) \sim \tfrac{1}{2}\varepsilon_d\nu_d, \quad d \to \infty. \tag{57}$$

Also, note that $\varepsilon_d^2\nu_d \to \infty$ by assumption. Hence, the first term in (55) is $\sim -\log(\varepsilon_d^2\nu_d)$ while the second term (including sign) is $\sim -\frac{1}{4}\varepsilon_d^2\nu_d$, so that the sum tends to $-\infty$.

Right tail. The treatment of the right tail is analogous, starting with (40) instead of (23). Details are omitted. □

In the following we examine the behavior of the normalizing constant c_d as $d \to \infty$.

Proposition 2. *In the setting considered here,*

$$\frac{1}{c_d} \sim \sqrt{\frac{2\pi}{\nu_d}}\, u_d^{d+1}\psi(u_d), \quad d \to \infty. \tag{58}$$

Proof. Let ε_d be such that $1/\nu_d \ll \varepsilon_d^2 \ll \omega_d/\nu_d$. Applying Proposition 1 and then performing a change of variables, we get

$$\frac{1}{c_d} = \int_0^\infty u^d\psi(u)\mathrm{d}u \sim \int_{(1-\varepsilon_d)u_d}^{(1+\varepsilon_d)u_d} u^d\psi(u)\mathrm{d}u \tag{59}$$

$$\sim \int_{-\varepsilon_d}^{\varepsilon_d} [(1+t)u_d]^d\psi[(1+t)u_d]u_d\mathrm{d}t \tag{60}$$

$$= u_d^{d+1}\psi(u_d)\int_{-\varepsilon_d}^{\varepsilon_d}(1+t)^d\frac{\psi[(1+t)u_d]}{\psi(u_d)}\mathrm{d}t. \tag{61}$$

As before,

$$\log\Big\{(1+t)^d\frac{\psi[(1+t)u_d]}{\psi(u_d)}\Big\} = d\log(1+t) - \Lambda[(1+t)u_d] + \Lambda(u_d) \tag{62}$$

$$= d\int_0^t \frac{s}{1+s}\mathrm{d}s - \int_0^t u_d\Lambda'[(1+s)u_d]\mathrm{d}s \tag{63}$$

$$= \int_0^t \frac{1}{1+s}\Big\{d - L[(1+s)u_d]\Big\}\mathrm{d}s. \tag{64}$$

Noting that $|s| \le \varepsilon_d$, and using (48), we get

$$-s\nu_d - s\nu_d/\omega_d \le d - L[(1+s)u_d] = L(u_d) - L[(1+s)u_d] \le -s\nu_d + s\nu_d/\omega_d. \tag{65}$$

Hence,

$$-\nu_d(1+1/\omega_d)\int_0^t \frac{s}{1+s}\mathrm{d}s \le \int_0^t \frac{1}{1+s}\Big\{d - L[(1+s)u_d]\Big\}\mathrm{d}s \le -\nu_d(1-1/\omega_d)\int_0^t \frac{s}{1+s}\mathrm{d}s, \tag{66}$$

with

$$\int_0^t \frac{s}{1+s}\mathrm{d}s = \tfrac{1}{2}t^2 + O(t^3) = \tfrac{1}{2}t^2 + O(\varepsilon_d^3), \tag{67}$$

since $|t| \le \varepsilon_d$, so that

$$\int_0^t \frac{1}{1+s}\Big\{d - L[(1+s)u_d]\Big\}\mathrm{d}s = -\tfrac{1}{2}t^2\nu_d + O(\varepsilon_d + 1/\omega_d)\varepsilon_d^2\nu_d, \tag{68}$$

where the big-O is uniform in $t \in [-\varepsilon_d, \varepsilon_d]$. We already took ε_d such that $(1/\omega_d)\varepsilon_d^2 \nu_d \to 0$, and it is compatible to choose ε_d such that, in addition, $\varepsilon_d^3 \nu_d \to 0$. When we do so, the remainder term above is $o(1)$, and in particular,

$$\int_0^t \frac{1}{1+s}\Big\{d - L[(1+s)u_d]\Big\}\mathrm{d}s = -\tfrac{1}{2}t^2\nu_d + o(1), \tag{69}$$

where the $o(1)$ term is uniform in $t \in [-\varepsilon_d, \varepsilon_d]$. With such a choice of ε_d, we continue our derivations above

$$\int_{-\varepsilon_d}^{\varepsilon_d}(1+t)^d\frac{\psi[(1+t)u_d]}{\psi(u_d)}\mathrm{d}t = \int_{-\varepsilon_d}^{\varepsilon_d}\exp\Big\{-\tfrac{1}{2}t^2\nu_d + o(1)\Big\}\mathrm{d}t \tag{70}$$

$$\sim \frac{1}{\sqrt{\nu_d}}\int_{-\varepsilon_d\sqrt{\nu_d}}^{\varepsilon_d\sqrt{\nu_d}}\exp\Big\{-\tfrac{1}{2}s^2\Big\}\mathrm{d}s \sim \frac{1}{\sqrt{\nu_d}}\sqrt{2\pi}, \tag{71}$$

since $\varepsilon_d\sqrt{\nu_d} \to \infty$ by assumption. □

We are finally equipped to establish a convergence in distribution for U_d.

Theorem 4. *In the setting considered here, $\sqrt{\nu_d}(U_d/u_d - 1)$ converges weakly to the standard normal distribution as $d \to \infty$.*

Example 1 (Continued). It can be checked that the same example of shape function ψ satisfies the conditions assumed here, with $uL'(u) \sim c\beta^2(\log u)^\alpha u^\beta$ as $u \to \infty$, so that

$$\nu_d = u_d L'(u_d) \sim \beta d, \quad d \to \infty. \tag{72}$$

Proof. Fix $r \in \mathbb{R}$ and let ε_d be as before. As in (61), we get

$$\mathbb{P}(\sqrt{\nu_d}(U_d/u_d - 1) \le r) = \mathbb{P}(U_d \le (1 + r/\sqrt{\nu_d})u_d) \tag{73}$$

$$\sim \mathbb{P}((1-\varepsilon_d)u_d \le U_d \le (1 + r/\sqrt{\nu_d})u_d) \tag{74}$$

$$= c_d u_d^{d+1}\psi(u_d)\int_{-\varepsilon_d}^{r/\sqrt{\nu_d}}(1+t)^d\frac{\psi[(1+t)u_d]}{\psi(u_d)}\mathrm{d}t. \tag{75}$$

Again, as before,

$$\int_{-\varepsilon_d}^{r/\sqrt{\nu_d}}(1+t)^d\frac{\psi[(1+t)u_d]}{\psi(u_d)}\mathrm{d}t \sim \int_{-\varepsilon_d}^{r/\sqrt{\nu_d}}\exp\Big\{-\tfrac{1}{2}t^2\nu_d\Big\}\mathrm{d}t \tag{76}$$

$$= \frac{1}{\sqrt{\nu_d}}\int_{-\varepsilon_d\sqrt{\nu_d}}^{r}\exp\Big\{-\tfrac{1}{2}s^2\Big\}\mathrm{d}s \sim \frac{1}{\sqrt{\nu_d}}\sqrt{2\pi}\Phi(r), \tag{77}$$

where Φ is the standard normal distribution function. We then combine this with Proposition 2. □

4 Discussion

While there is relatively little related work, a detailed comparison with (Sherlock and Elton, 2012) is in order.[1] Sherlock and Elton focus on the non-compact case — corresponding to Section 3 here. They derive the same result as our Theorem 3 under different conditions. They require that $\eta(u) := \Lambda(\exp(u))$ is twice differentiable with $\eta''(u) \to \infty$ as $u \to \infty$, while our condition is a bit

[1] Note that we only became aware of that work after we posted our paper on `arxiv.org`, which explains some of the nontrivial overlap.

weaker than requiring that η is once differentiable with $\eta'(u)/u \to \infty$ and increasing. Note that their condition is equivalent to requiring that L is differentiable with $uL'(u) \to \infty$, a condition that arises in Section 3.2. Sherlock and Elton do not establish weak convergence, however, but they obtain other results. In particular, they establish concentration for a marginal of X_d and also for the maximum of the marginals. In addition, they extend their results to the case of elliptical distributions under conditions on the eigenvalues of the scaling matrix.

What are possible consequences for statistical modeling? Because of the weak convergence of the sort established here, the behavior of U_d is asymptotically characterized solely by a few parameters of the underlying distribution. For example, if the conditions of Theorem 2 are fulfilled, then the distribution of U_d in the large-dimension limit ($d \to \infty$) only depends on u_* (irrelevant in practice because scale is typically estimated) and the behavior of ψ near u_*. In particular, whether $\psi(u) = \mathbb{I}\{u \le 1\}$ or $\psi(u) = (2-u)\mathbb{I}\{u \le 1\}$, in both cases, $d(1-U_d)$ converges weakly to the exponential distribution with rate 1. This means that, in order to even distinguish two such distributions with nontrivial accuracy, we require a sample of size that increases with d. (We did not attempt to quantify this further, although this is possible by framing the problem as a hypothesis testing problem.) A similar phenomenon arises with certain distributions with non-compact support, based on our Theorem 4. Thus, if the sample size is small relative to the dimension, very different shape functions (meaning, different ψ's) could yield indistinguishable models.

The flip side of this is a form of universality of the Gaussian distribution, in particular, in context such as Linear Discriminant Analysis (classification) or Gaussian Mixture Modeling (clustering). Surely, both LDA and GMM have computational advantages over other methods (the latter using the EM algorithm, for example). Beyond this important computational aspect, our results indicate that when the sample is small relative to the dimension, fitting a Gaussian model may be, in fact, indistinguishable from fitting another model base on a shape function having similar characteristics as the standard normal distribution that dictate the asymptotic behavior of U_d.

References

Anderson, T. W. (2003). *An introduction to multivariate statistical analysis* (Third ed.). Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ.

Bhattacharyya, S. and P. J. Bickel (2015). Adaptive estimation in elliptical distributions with extensions to high dimensions. Technical report, University of California, Berkeley.

Bickel, P. J., C. A. Klaassen, Y. Ritov, J. A. Wellner, et al. (1998). Efficient and adaptive estimation for semiparametric models.

Boucheron, S., G. Lugosi, and P. Massart (2013). *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press.

Chang, G. T. and G. Walther (2007). Clustering with mixtures of log-concave distributions. *Computational Statistics & Data Analysis 51*(12), 6242–6251.

Diaconis, P. and D. Freedman (1984). Asymptotics of graphical projection pursuit. *The Annals of Statistics 12*(3), 793–815.

Diaconis, P. and D. Freedman (1987). A dozen de finetti-style results in search of a theory. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques 23*(S2), 397–423.

Ledoux, M. (2005). *The concentration of measure phenomenon*. Number 89. American Mathematical Soc.

Lewis, D. D. (1998). Naive (bayes) at forty: The independence assumption in information retrieval. In *European conference on machine learning*, pp. 4–15. Springer.

Sherlock, C. and D. Elton (2012). A class of spherical and elliptical distributions with gaussian-like limit properties. *Journal of Probability and Statistics 2012*.